Name ______________________ Student No. ______________________

No aids allowed. Answer all questions on test paper. Use backs of sheets if necessary.

Total Marks: **30**

[10] 1. Consider the figure below. Assuming TCP Reno is the protocol experiencing the behavior shown in the figure, answer the following questions (on the next page). In all cases you should include a *brief* discussion justifying your answer; when interpreting the diagram, approximate the numbers visually — the grader will allow for some "wiggle room" when writing down the values.

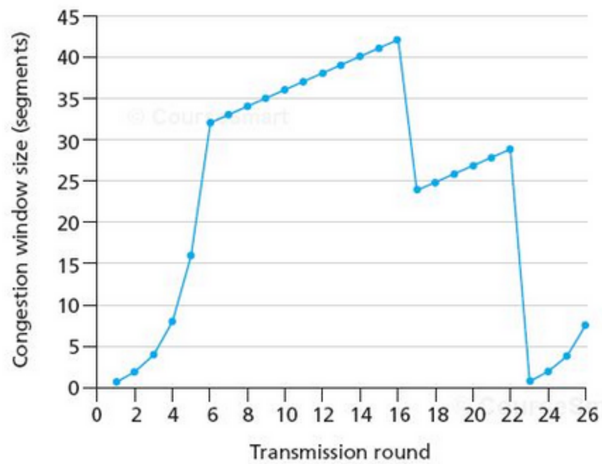

(a) Identify the intervals of time when TCP slow start is operating.

(b) Identify the intervals of time when TCP congestion avoidance is operating.

(c) After the 16th transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

(d) After the 22nd transmission round, is segment loss detected by a triple duplicate ACK or by a timeout?

(e) What is the initial value of `ssthresh` at the first transmission round?

(f) What is the value of `ssthresh` at the 18th transmission round?

(g) What is the value of `ssthresh` at the 24th transmission round?

(h) During what transmission round is the 70th segment sent?

(i) Assuming a packet loss is detected after the 26th round by the receipt of a triple duplicate ACK, what will be the values of the congestion window size of `ssthresh`?

Empty page for answers to question 1.

Solution:

(a) TCP slowstart is operating in the intervals $[1, 6]$ and $[23, 26]$.

(b) TCP congestion avoidance is operating in the intervals $[6, 16]$ and $[17, 22]$.

(c) After the 16th transmission round, packet loss is recognized by a triple duplicate ACK. If there was a timeout, the congestion window size would have dropped to 1.

(d) After the 22nd transmission round, segment loss is detected due to timeout, and hence the congestion window size is set to 1.

(e) The threshold is initially 32, since it is at this window size that slow start stops and congestion avoidance begins.

(f) The threshold is set to half the value of the congestion window when packet loss is detected. When loss is detected during transmission round 16, the congestion windows size is 42. Hence the threshold is 21 during the 18th transmission round.

(g) The threshold is set to half the value of the congestion window when packet loss is detected. When loss is detected during transmission round 22, the congestion windows size is 26. Hence the threshold is 13 during the 24th transmission round.

(h) During the 1st transmission round, packet 1 is sent; packet 2-3 are sent in the 2nd transmission round; packets 4-7 are sent in the 3rd transmission round; packets 8-15 are sent in the 4th transmission round; packets 16-31 are sent in the 5th transmission round; packets 32-63 are sent in the 6th transmission round; packets 64-96 are sent in the 7th transmission round. Thus packet 70 is sent in the 7th transmission round.

(i) The threshold will be set to half the current value of the congestion window (8) when the loss occurred and congestion window will be set to the new threshold value $+3 \cdot \text{MSS}$. Thus the new values of the threshold and window will be 4 and 7 respectively (assuming MSS $=$ 1, or, it would also be correct to say cwnd $=$ $4 + 3 \cdot \text{MSS}$).

[10] 2. Consider the network shown below. Suppose AS3 and AS2 are running OSPF for their intra-AS routing protocol. Suppose AS1 and AS4 are running RIP for their intra-AS routing protocol. Suppose eBGP and iBGP are used for the inter-AS routing protocol.

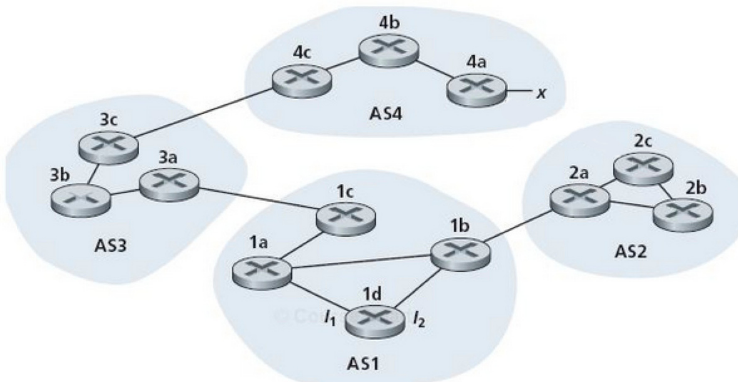

(a) Router 3c learns about prefix x from which routing protocol: OSPF, RIP, eBGP, or iBGP?

(b) Router 3a learns about x from which routing protocol?

(c) Router 1c learns about x from which routing protocol?

(d) Router 1d learns about x from which routing protocol?

Solution: (a) eBGP; (b) iBGP; (c) eBGP; (d) iBGP
The reason is simply this (see pg. 393 in Kurose, 6th edition): a BGP session that spans two ASs is called an *external BGP* (eBGP), while a BGP session between routers in the same AS is called an *internal BGP* (iBGP).

[10] 3. Consider the subnets shown below. Suppose that we replace the router between subnet 1 and subnet 2 with a switch S1, and label the router between subnet 2 and subnet 3 as R1.

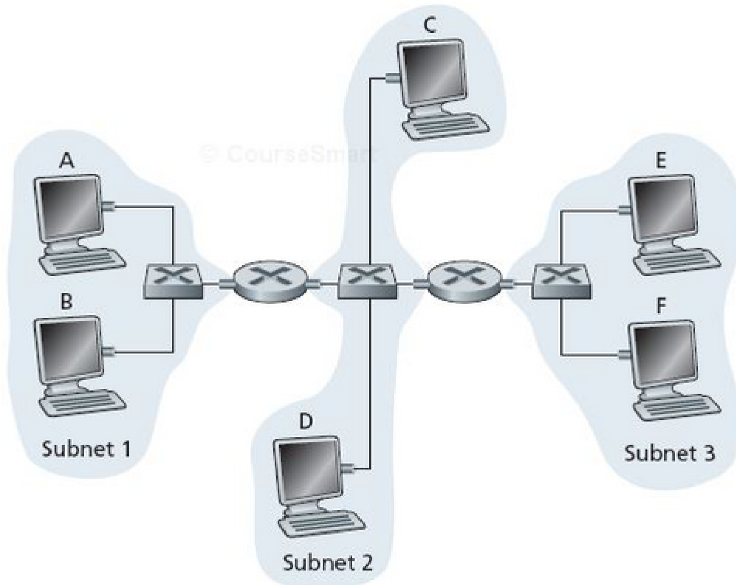

(a) Consider sending an IP datagram from Host E to Host F. Will Host E ask router R1 to help forward the datagram? Why? In the Ethernet frame containing the IP datagram, what are the source and destination IP and MAC addresses? (You can answer this question by saying "Host X's IP address," etc.)

(b) Suppose E would like to send an IP datagram to B, and assume that E's ARP cache does not contain B's MAC address. Will E perform an ARP query to find B's MAC address? Why? In the Ethernet frame (containing the IP datagram destined to B) that is delivered to router R1, what are the source and destination IP and MAC addresses?

(c) Suppose Host A would like to send an IP datagram to Host B, and neither A's ARP cache contains B's MAC address nor does B's ARP cache contain A's MAC address. Further, suppose that the switch S1's forwarding table contains entries for Host B and router R1 only. Thus, A will broadcast an ARP request message. What actions will switch S1 perform once it receives the ARP request message? Will router R1 also receive this ARP request message? If so, will R1 forward the message to Subnet 3? Once Host B receives this ARP request message, it will send back to Host A an ARP response message. But will it send an ARP query message to ask for A's MAC address? Why? What will switch S1 do once it receives an ARP response message from Host B?

Empty page for answers to question 3.

Solution:

(a) No. E can check the subnet prefix of Host F's IP address, and then learn that F is on the same LAN. Thus, E will not send the packet to the default router R1.

Ethernet frame from E to F:

Source IP = E's IP address
Destination IP = F's IP address
Source MAC = E's MAC address
Destination MAC = F's MAC address

(b) No, because they are not on the same LAN. E can find out that they are not on the same subnet by checking B's IP address.

Ethernet frame from E to R1:

Source IP = E's IP address
Destination IP = B's IP address
Source MAC = E's MAC address
Destination MAC = The MAC address of R1's interface connecting to Subnet 3.

(c) Switch S1 will broadcast the Ethernet frame on the interface to subnet 2 (since it got the frame on its interface to subnet 1) as the received ARP frame's destination address is a broadcast address. And it learns that A resides on Subnet 1 which is connected to S1 at the interface connecting to Subnet 1. And, S1 will update its forwarding table to include an entry for Host A.

Yes, router R1 also receives this ARP request message, but R1 won't forward the message to Subnet 3.

B won't send ARP query message asking for A's MAC address, as this address can be obtained from A's query message.

Once switch S1 receives B's response message, it will not add an entry for host B to its forwarding table (as the question says explicitly that S1 already had an entry for host B), and then drop the received frame as destination host A is on the same interface as host B (i.e., A and B are on the same LAN segment).